arXiv:1405.5732v2 [cs.CV] 26 May 2014

Self-tuned Visual Subclass Learning with Shared Samples An Incremental Approach

Updated ICCV 2013 Submission

Hossein Azizpour
Royal Insitute of Technology(KTH)
azizpour@kth.se

Stefan Carlsson
Royal Insitute of Technology(KTH)
stefanc@csc.kth.se

Abstract

Computer vision tasks are traditionally defined and evaluated using semantic categories. However, it is known to the field that semantic classes do not necessarily correspond to a unique visual class (e.g. inside and outside of a car). Furthermore, many of the feasible learning techniques at hand cannot model a visual class which appears consistent to the human eye. These problems have motivated the use of 1) Unsupervised or supervised clustering as a preprocessing step to identify the visual subclasses to be used in a mixture-of-experts learning regime. 2) Felzenszwalb et al. part model and other works model mixture assignment with latent variables which is optimized during learning 3) Highly non-linear classifiers which are inherently capable of modelling multi-modal input space but are inefficient at the test time. In this work, we promote an incremental view over the recognition of semantic classes with varied appearances. We propose an optimization technique which incrementally finds maximal visual subclasses in a regularized risk minimization framework. Our proposed approach unifies the clustering and classification steps in a single algorithm. The importance of this approach is its compliance with the classification via the fact that it does not need to know about the number of clusters, the representation and similarity measures used in pre-processing clustering methods a priori. Following this approach we show both qualitatively and quantitatively significant results. We show that the visual subclasses demonstrate a long tail distribution. Finally, we show that state of the art object detection methods (e.g. DPM) are unable to use the tails of this distribution comprising 50% of the training samples. In fact we show that DPM performance slightly increases on average by the removal of this half of the data.

1. Introduction

A limited set of visual classes such as frontal faces, pedestrians and other objects detected in the context of urban traffic admit a rate of detection that challenges that of human vision and make them interesting for commercial applications([24, 25]). This can be attributed to the relatively limited degree of intra class variation expressed by the images in these classes . On the other hand, multi view detection of general object classes show a performance in the range of around 20 to 65 % in average precision (AP) ([10]). In recent years, increase in the performance has come about by various algorithms from year to year. They are designed based on a limited set of training data and are expected to capture the class variation which is substantial due to different camera viewpoint, subordinate classes, etc. So, it is still an open question whether this approach will eventually lead to the level of the performance of human visual system. The complexity of the classifiers to be learnt for these datasets is simply too large for the limited sets of data used to train them. The obvious remedy of trying to increase training data does not necessarily lead to better classifiers due to the emerging influence of outliers([35]). To get around this we suggest an alternative way of designing visual classifiers. We consider the ultimate goal of a classifier is to achieve performance at the level of human vision. However, instead of directly designing a classifier that works for a full training set, we examine subsets of the class for which classifiers can be learnt with good performance. The objective of this design is to increase the size of these subsets and possibly merge them while maintaining the good classification rate. The ”size” has to relate in some sense to the variation within the subclass. As these subclasses grow and merge with each other it is required to increase the complexity of the classifiers applied to them. This approach should be contrasted with the prevailing methods of designing classifiers for the full dataset that usually lead to complex non-linear optimization problems. In order to handle these, the optimization methods have to

Figure 1: A subset of visual subcategories found by our framework for the class of bicycle in the PASCAL VOC 2007 training set. The subclasses can share samples as long as they keep high accuracy on predicting them. The corresponding HOG filters for each subclass are marked by the same color of the border.[1]

be regularized in order to avoid over fitting. The proper information for regularization is simply not there at the start but has to be acquired gradually from the data itself. We further claim that instead of directly competing in the performance of the final classification rate, competition should be in terms of size and complexity of subclasses that are properly classified.

In the current work, we devise an incremental discriminative procedure which is able to simultaneously discover and learn such visual subclasses (Figure 1). In particular, we formulate the learning problem as regularized risk minimization to find the largest subset which can be reliably detected. Taking the subset out we continue the same procedure on the remaining samples and so forth. Intuitively, this will leave us with a long-tailed distribution of subclasses (Figure 2). For the final experiments we specialize to the deformable part models (DPM) [11] as our RRM framework.
This iterative procedure helps to keep outliers –that might harm the performance of subclasses– away as isolated exemplars. This means that our work relates primarily to the so called mixture models of classification [11] [32] [1]. This incremental approach, by finding "visual subclasses", allows understanding of the problem at hand by analyzing the core large subsets *and* isolated exemplars. It should be mentioned that by a "visual subclass" we mean a subset of our semantically defined class that can be reliably detected given a learning framework and feature representation and thus does not necessarily comply with visual subclasses to the human eye. Saying that, however, we can see in Figure (1) that the subclasses are still meaningful in appearance. Finally, we posit that the recent success of Deep Learning methods trained on large outside-training datasets for visual recognition tasks [27] is related to unification of learning representations that are invariant to such subclass variations and classification. The recognition problem is still unsolved, thus the advantage of our procedure to the deep learning methods is that we have a systematic analysis of where our recognition system fails which similarly highlights the importance of handling the tails of the distribution in deep learning frameworks.

1.1. Contributions and Outline

The contributions of this work can be listed as follows,
I) A general incremental view on the recognition task is motivated: i.e. the identification, growth and merging of the *maximal* subclasses which can be detected with *high accuracy*.
II) We propose a simple yet effective discriminative procedure based on self-paced learning formulation for simultaneous discovery and learning of such subclasses (without the prior knowledge about the number of subclasses, the representation space or similarity measure suitable for clustering) which fits into the regularized risk minimization scenario. The formulation, seen as a data-driven clustering framework, is able to span the spectrum between the two extremes of a)one exemplar classifier per sample and b)a monolithic classifier for all samples.
III) We show the proposed self-paced learning is able to detect outliers (w.r.t. to the base classifier) and is equivalent to the ramp loss formulation of SVM.
IV) We allow sharing of the samples between the mixtures

[1]This figure is generated automatically using a subset of each subclass members, for larger pictures of all the classes please refer to the supplementary material.

and we show that this helps the final classification by regularizing the decision boundaries where different mixture classifiers meet.
V) We show that the visual subclasses follows a long tail distribution.
VI) By an aggregation of the identified *non-singular* subclasses (covering 15 to 80% of the data) we show that state of the art methods (e.g. DPM) are unable to use, on average, 50% of data appearing at the tails of this distribution.
VII) We further investigate the statistics of the isolated exemplars in terms of different image properties which gives an insight for future works.

In the next section we discuss the related works and highlight the differences. It is followed by the explanation of our method for learning visual subclasses. Then we report and discuss results from a set of experiments supporting the model aspects. Finally, we conclude the paper and mention interesting future directions.

2. Related Works

The idea of gradually designing learners is at the heart of the methodology known as incremental learning [33, 14, 20]. However unlike standard incremental learning we focus on the gradual construction of well classified subsets and we assume that all training data is available at the start, although incrementally adding training data could be part of our scheme too.
Identification of the regions of the input space to be classified by each expert is the main question when using mixture of experts. Particularly for the task of classification on the more realistic vision datasets[10], different works exist. They are mainly comprised of partitioning the input space in some way [31] or hard disjoint partitioning of the input samples [6] and then learning individual classifiers using those subsets. Some methods are provided with explicit assignment cluster members. [17] uses viewpoint annotations. However, annotation of other sources of variation, e.g. subordinates is hard to define and costly to do. Other methods, starting with an initial clustering, employ an EM-like approach to update the clusters. This is done by letting different experts compete for samples [11, 21, 28]. These methods, however, are susceptible to initialization.
To improve the initialization, other works use a clustering methods on either the image patches itself [1] or on some available meta-data [4, 2, 16]. While they have indeed shown improvements over baselines, the number of clusters is unknown a priori. Furthermore, the clustering method might not be compatible with the final model since they are optimized for different objectives.
Our formulation is similar to that of self-paced learning [22], however, with a completely different perspective. Kumar et al. use the formulation to cope with the sensitiveness of optimizing latent variable models. And for this they gradually consider all the examples by starting from simpler examples. This is in contrast to our scheme. We want each subset to contain only the accurately predictable examples. So, conversely, we start by including more examples at the start and gradually converge to the largest core of the samples that can be reliably detected.
Our framework, can be seen as a discriminative clustering. Our approach, however, is different from other existing discriminative clustering methods [3, 13, 18] in the sense that it tries to best discriminate the positive and negative classes. This is unlike the conventional discrimination of the positive clusters that we actually avoid by sharing samples. [26, 29] also demonstrate the effectiveness of overlapping clusters. In addition our proposed clustering method does not require prior knowledge about the number of subclasses, the representation space or similarity measure suitable for clustering.
Finally, our framework is closely related to [8, 9] for discovery of midlevel features used in tasks other than detection. However, we see our proposed method as a more principled strategy for discovery of subclasses which can be used to group such mid-level representations as well.

3. Methodology

Here we explain the framework used to train our self-tuned mixture classifiers. First we recap the Regularized Risk Minimization problem (RRM) and mixture of experts(MEX). Then we introduce a new formulation for learning each expert of MEX and discuss its properties. Finally, by injecting DPM as the individual RRM learner we show how we optimize the objective in an iterative manner.

3.1. Mixture of Regularized Risk Minimizers

Discriminative learning methods have been very successful in many applications of machine learning such as computer vision. This is particularly true for classification tasks when provided with a relatively large set of labelled training samples. In the context of *discriminative* learning of binary classifiers, given a set of N labelled training samples $D = \{\langle x_i, y_i \rangle\}_{i=1}^{N}, y_i \in \{-1, 1\}$, we are usually faced with a RRM formulation of the form,

$$\min_{\mathbf{w}} J(\mathbf{w}) := \Omega(\mathbf{w}) + \lambda R_{emp}(\mathbf{w}, D) \tag{1}$$

$$R_{emp}(\mathbf{w}, D) = \sum_{i=1}^{N} l_i^{\mathbf{w}} \tag{2}$$

where $l_i^{\mathbf{w}}$ is the "loss" of sample i. Solving an RRM involves a trade-off between the regularization term Ω and the risk term R_{emp} which is controlled by the coefficient λ. There are different choices of a reguralizer e.g. L1, L2

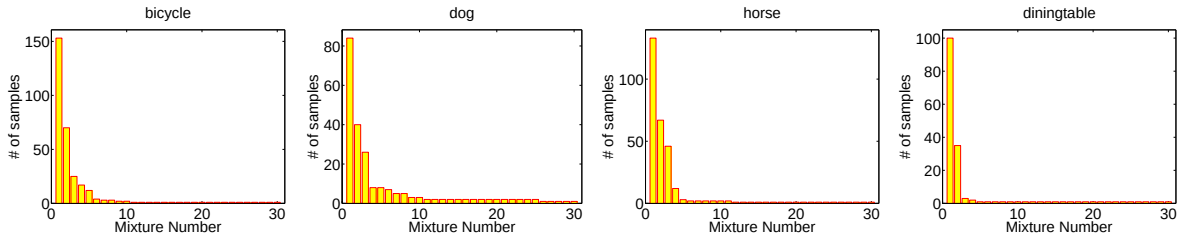

Figure 2: Distribution of the number of samples inside each visual subclass sorted in descending order for the first 30 components. A long tailed distribution is observed.

and loss functions e.g. least squares or hinge (for Lasso or Support Vector Machines respectively). Depending on these choices there are well-established solvers in the literature.

However, it has been observed in several applications that the distribution of the training samples in the input space is multi-modal. This is because of the semantically defined labels, data incompleteness and poor representation (e.g.. not invariant enough to different conditions). Although kernel methods or similar approaches can partially overcome these problems, they are usually inefficient for large scale problems and parameters should be carefully tuned for them to avoid overfitting. This has caused the mixture-of-experts (MEX) and ensemble methods to become popular. MEX uses a combination of simpler learners (e.g. linear SVMs) to train a more complicated classifier able to cope better with the multi-modality of the input space. It should be noted that MEX is subtly different from other ensemble methods in that it divide the input space by training each expert on either a hard subset or non-uniform distribution of training data. MEX of M learners in terms of RRMs is defined as follows,

$$\min_{W} J(W) := \Omega(W) + \lambda R(W, D) \tag{3}$$

$$R(W, D) = \sum_{i=1}^{N} g(l_i^{\mathbf{w}_1}, l_i^{\mathbf{w}_2}, ..., l_i^{\mathbf{w}_M}) \tag{4}$$

$$W = \{\mathbf{w}_1, \mathbf{w}_2, ..., \mathbf{w}_M\} \tag{5}$$

where $g(.)$ is an aggregation function which can be anything from a gating function or a minimum of the losses produced by different experts. The learning of MEX involves two main steps, 1) identification and Learning of the parameters of the individual experts, 2) learning the parameters of the aggregation function. Some methods combine these two steps though [11]. Due to this design, the three major difficulties of training a MEX are first, how to identify the partitions of the input space to be learned by an expert, secondly MEX is very likely to have problems at the boundaries where expert learners meet (these boundary regions usually have low or zero sample support), and finally what is a good choice of an aggregation function g. In the next section we introduce a new incremental approach as an effort to address the first two problems. The third issue of choosing an aggregation function, however, remains open for future work.

3.2. Self-tuned Expert Learning

We approach the problem of visual subclass discovery in an incremental fashion. By that we mean we are interested in finding the *largest* subset of our positive training data ($y = 1$) which we can distinguish from the negative data with a high precision. Then taking that subset away we continue by finding the next largest subset and so on. This will help the classifiers to train on more well aligned samples due to the removal of outliers. This helps us to better understand the classification problem we face by looking at more isolated components. We pose the problem of finding the largest subset with high detection accuracy at each step using the following RRM optimization,[2]

$$\min_{\mathbf{w},\mathbf{v}} J'(\mathbf{w}) := \Omega(\mathbf{w}) + \lambda R'(\mathbf{w}, \mathbf{v}, D) \tag{6}$$

$$R'(\mathbf{w}, \mathbf{v}, D) = \sum_{i=1}^{N_p} v_i(l_i^{\mathbf{w}} - 1/K) + \sum_{i=1}^{N_n} l_{i+N_p}^{\mathbf{w}} \tag{7}$$
$$\mathbf{v} = \{v_1, v_2, ..., v_{N_p}\}, v_i \in \{0, 1\}$$

where $\mathbf{v}$ is a binary indicator vector that denotes which samples are included in the current mixture. K is a scalar which controls how predictable a sample should be to get in the subset. Large K is a tighter constraint on the predictability of the samples to be included. This is because as K gets larger the loss l_i of sample i should be lower to be included. This formulation is interestingly similar to that of Kumar et al's self-paced learning[22]. As described in [22] for a convex choice of l and Ω the relaxed formulation (i.e. $v_i \in [0, 1]$) is a biconvex optimization in $\mathbf{w}$ and $\mathbf{v}$. This means that when we fix either of the subset of variables $\mathbf{w}$

[2]For the simplicity of the notation, we assume samples are ordered by their labels thus in the set D there are N_p positive ($y_i = 1$) samples indexed by $i = 1 : N_p$ and then N_n negative ($y_i = -1$) samples indexed by $i = N_p + 1 : N_p + N_n$.

or $\mathbf{v}$ the problem becomes convex over the other subset. Biconvex problems are specially interesting since they have good approximate solvers [15] as well as branch and bound global optimizers.

Note that by choosing large enough or small enough K we span a spectrum where we have Exemplar learners[23] at one extreme and one single learner at the other respectively. It should be further noted that the formulation of self-paced learning is equivalent to the ramp loss formulation of [7]. Ramp loss was originally designed to reduce the number of support vectors by ignoring the outliers which is essentially what we want in this approach. Further investigation of this equivalence and its consequences appears at appendix A in Sec 6.

Decision at boundaries where two independently trained experts meet can become ambiguous. To deal with this issue, after discovering each subclass, we let it share the already selected samples of previously trained experts. In other words we fix the corresponding elements of $\mathbf{v}$ to 1 for the discovered subset and optimize (11) for the whole training set. This will allow for sharing of samples between the experts which makes them become smoother at the boundaries and potentially fill the gaps in the input space. In the experiments we found this is a crucial step for the performance of each expert.

Finally, this framework can be seen as a way of *discriminative* incremental clustering. In that sense this clustering can potentially perform superior to conventional pre-processing clustering methods used for MEX. This is due to the fact that it unifies the partitioning of the positive samples in a way to comply better with the final classifier. That is our procedure does not require the number of subclasses, a new representation space or similarity measure suitable for clustering.

3.3. Self-tuned DPM

To evaluate the effectiveness of our method for visual recognition we test it for the task of object detection. Since deformable part models (DPM) of Felzenszwalb at. al.[11] is the current state-of-the-art, we use it as our individual experts. The formulation by replacing RRM terms with those of DPM then becomes as follows,

$$\min_{\mathbf{w},\mathbf{v}} J'(\mathbf{w},\mathbf{v}) := \|\mathbf{w}\|^2 + \lambda(\sum_{i=1}^{N_p} v_i(\max(0, 1 - f_{\mathbf{w}}(x_i)) - 1/K) + \sum_{i=1}^{N_n} \max(0, 1 + f_{\mathbf{w}}(x_{i+N_p}))) \quad (8)$$
$$\mathbf{v} = \{v_1, v_2, ..., v_{N_p}\}, v_i \in \{0, 1\}$$

$$f_{\mathbf{w}}(x_i) = \max_{h_i}(\mathbf{w}^T \Phi(x_i, h_i)) \quad (9)$$

where $\Phi(.,.)$ is the HOG feature vector extracted from image x_i at the location indicated by part latent variables h_i. DPM uses latent variables for modelling part positions in the scoring function f; Note that the loss function of positives in Eq. 8 is a concave function. Thus (8) is not a biconvex optimization anymore. However, as illustrated in [11] the problems become convex as soon as one fixes the latent variables for the positive samples. We use this in the next section to come up with an Alternative Convex Search (ACS) algorithm to solve the above optimization problem.

3.4. Optimization

Solution of biconvex optimization problems can be well approximated using Alternative Convex Search (ACS) approach [15]. This has been confirmed for Kumar et al.'s self-paced learning formulation of the structured-SVM [22]. Using ACS one should fix one of the two disjoint subset of variables and solve the induced convex problem for the other subset and vice versa, and alternate until no improvement is possible. On the other hand Latent SVM formulation of DPM can be solved by alternative search between fixing latent variables and optimizing for $\mathbf{w}$ in the induced objective function. When $\mathbf{w}$ is fixed, the optimal solution of v and $f_{\mathbf{w}}$ can be found in a single step (see Alg. 1).

Data: D, K, C_p, C_n
return $\mathbf{w}$ *model and* $\mathbf{v}$ *indicator vector*
$\mathbf{v} \leftarrow \{1\}$
while $\mathbf{v} \neq \mathbf{v}_{old}$ **do**
  0: $\mathbf{v}_{old} \leftarrow \mathbf{v}$
  1: Solve for $\mathbf{w}$ the convex program (11) with fixed $\mathbf{v}$
  2: **for** $i = 1 : N_p$ **do**
    if $1 - \mathbf{w}^T \mathbf{x}_i < 1/K$ **then**
      $v_i = 0$
    end
    else
      $v_i = 1$
    end
  end
end

Algorithm 1: ACS for self-paced learning

In an incremental approach, as illustrated in Alg. 2 we solve the equation (8) initially with the whole training set with Alg. (1) and then taking the indicated samples by $\mathbf{v}$ away, we repeat the procedure until no sample remains. It should be noted since we have an object detection scenario, we have huge number of possible negative windows from the negative images. Therefore, we adopt the hard negative mining approach of [11] at step 4 of the algorithm (1).

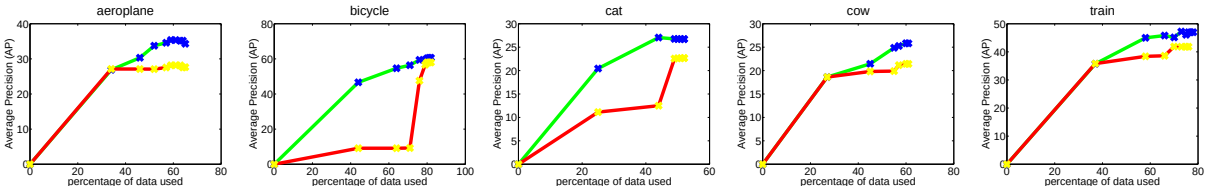

Figure 3: The evoultion of the AP when adding more experts for our full system (green -light gray- with blue crosses) compared with the case without sharing samples (red -dark gray- with yellow crosses) on 5 classes of VOC07.

Data: D, K, C_p, C_n
return *W set of trained experts V set of indicator vectors*
while $D' \neq \emptyset$ **do**
  0: $D' \leftarrow D$
    $W, V \leftarrow \emptyset$
  1: Solve (8) for $\mathbf{v}$ and $\mathbf{w}$ using alg.1
  2: $V \leftarrow V \cup \mathbf{v}$
  3: $W \leftarrow W \cup \mathbf{w}$
  4: $D' \leftarrow D' \backslash \{i | v_i = 1\}$
end

Algorithm 2: Self-paced incremental

4. Experiments

We have chosen the task of Object Detection for evaluation of our method. The challenging dataset of PASCAL VOC 2007 is used as the testbed. In the following sections, we give the details of our implementation and then we continue by presenting the final detection performance of our system as well as showing quantitative and qualitative results illustrating different aspects of our framework.

4.1. Implementation Details

We use the release 5 of the DPM [12] as our base code without post processing of context re-scoring. We cross-validated a few choices of the parameters for the classes of "horse" and "car" on the validation set.
We set K so that it contains all the samples on the positive half space defined by the decision boundary $\mathbf{w}$ ($K = 0.7$). Furthermore, we found that starting with a lower K in ACS (Alg. 1) and then gradually increasing it to the final value will produce better results. It is initialized to a value that roughly includes the samples inside the SVM margin and on the positive half space(0.55). Then K is increased by a factor of 1.05 to reach its final value.
This helps the optimization to better find the core subset at each iteration by taking relatively more samples in the beginning and then decreasing them if necessary.
This approach is in contrast to [22] where they start with higher value of K and decrease it gradually. The aim there, is to consider all the samples eventually. They want to include cleaner sample first for the non-convex optimization of latent variables and then add the rest. We also tried to decrease the K from a higher value and found that increasing K gradually is a better choice for our purpose. The same gradual decrease in the λ, when the subset gets larger, is applied which in practice helped the convergence of the alternative search.
We do not have any particular aggregation policy to measure the final detection performance of the set of expert classifiers W. All the detections from all the experts are simply put together into the same pool (as if they were produced by a single classifier) and then a non-maximum suppression based on PASCAL VOC overlap criterion is applied.

4.2. Object Detection on PASCAL VOC 2007

To evaluate the final object detection performance, in favour of computation, we use a subset of trained experts. We only consider subsets that have at least 5 samples. This leaves us with less than 10 experts for each class. Table (1) shows the final Average Precision(AP) on the 20 classes of VOC07 and it is compared with many of the existing state-of-the-art methods. Our method uses half of the data in average and yet it performs comparably to all the state-of-the-art methods. We can see that there is a marginal improvement in terms of mean AP over all classes. It is worth of mentioning that [5], [30] and [16] use multiple features (in addition to HOG). [16] uses meta-data for clustering. Finally, we plot the evolution of the AP as we add more experts in Figure (8). Observe that most of the performance is achieved at the very early steps. The most interesting future work, indicated by these results, is how to benefit from the rest of the isolated samples to push for better performance.

4.3. Aggregation

To better understand the effect of a proper aggregation, we follow the work of Hoiem et al. in [19] and remove the false positives generated by poor localization. We remove the detections that overlap between 0.1 to 0.5 to the provided ground truth and re-compute the AP. We do the same procedure for our DPM baseline [12]. Table (2) shows that we have a larger margin in the mean AP and we outperform DPM in 16 out of 20 classes using the same fraction of data as mentioned in Table (1) for the "Ours" results.

	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
E-SVMs[23]	20.8	48.0	7.7	14.3	13.1	39.7	41.1	5.2	11.6	18.6	11.1	3.1	44.7	39.4	16.9	11.2	22.6	17.0	36.9	30.0	22.7
MCL[1]	29.0	50.1	9.6	15.0	18.9	41.1	49.7	10.3	16.0	21.0	17.0	10.3	50.0	39.6	33.0	9.0	19.8	22.0	38.2	34.3	26.7
[16]	33.4	37.0	15.0	15.0	22.6	43.1	49.3	32.8	11.5	**35.8**	17.8	16.3	43.6	38.2	29.8	11.6	**33.3**	23.5	30.2	39.6	29.0
AMH[5]	34.8	54.4	**15.5**	14.6	24.4	50.9	54.0	**33.5**	20.6	22.8	34.4	**24.1**	55.6	47.3	34.9	**18.1**	20.2	30.3	41.3	43.3	33.8
MKL[30]	**37.6**	47.8	15.3	15.3	21.9	50.7	50.6	30.0	17.3	33.0	22.5	21.5	51.2	45.5	23.3	12.4	23.9	28.5	45.3	**48.5**	32.1
DPM v5[12]	33.2	60.3	10.2	**16.1**	**27.3**	54.3	**58.2**	23.0	20.0	24.1	26.7	12.7	**58.1**	**48.2**	**43.2**	12.0	21.1	**36.1**	46.0	43.5	33.7
Ours	34.3	**60.7**	10.3	13.7	21.8	**57.5**	55.5	26.7	**22.0**	25.8	**38.3**	14.3	57.0	47.7	43.0	9.0	19.8	34.4	**47.0**	43.7	**34.2**
% of Data	65%	82%	28%	25%	45%	70%	73%	52%	43%	61%	66%	38%	74%	69%	63%	14%	36%	50%	78%	71%	55%

Table 1: Results of our method for object detection in VOC07 compared to other state-of-the-art baselines. The last row indicates the fraction of the training data used for each class to achieve the reported AP. It should be noted that [5],[30] and [16] use additional features to HOG. [16] uses meta-data for clustering.

	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
DPM v5[12]	37.3	62.2	10.4	**17.8**	**28.1**	56.8	**60.5**	27.9	22.0	24.8	30.0	13.4	**61.0**	49.5	52.4	14.3	21.9	38.2	51.7	44.3	36.2
Ours	**39.5**	**67.0**	**13.4**	16.4	22.7	**60.4**	58.5	**38.5**	**24.9**	**27.1**	**42.3**	**15.3**	59.6	**49.6**	**52.7**	**15.0**	**22.5**	**39.3**	**52.0**	**46.8**	**38.2**

Table 2: Average Precision results on VOC07 by taking out the localization false positives. As suggested by [19] we remove detections that have between 0.1 to 0.5 overlap with the corresponding ground truth. The results are by using the same fraction of the data as mentioned in Table (1)

This indicates the significantly higher ability of our proposed classifiers in recalling the objects and indicates the benefit of designing an appropriate localization scheme for aggregation of the different experts to reduce the false positives due to multiple counting. Furthermore, Fig. (8) shows that in most of the classes after addition of the 5th mixture the performance does not necessarily increase. Most of the performance is gained by addition of the first experts. This highlights the importance of score calibration as more mixtures are added [23].

4.4. Sharing Samples

On a subset of classes we show the evolution of the AP when adding more mixtures by comparing with the case when we do not allow sharing (Fig. 3). The gain in the improvement when sharing samples between mixtures indicate the importance of this procedure, possibly because it generates smoother boundaries and help the generalization of each of the experts for the final aggregation.

4.5. Qualitative Results

Fig. 4 depicts the HOG filters and a sample for the core subset of each of the classes. The sharp and clear edges exhibited in the filters show the good alignment of the samples caused by removal of outliers in the core subclasses.
To see figures illustrating the samples assigned to each expert and how they share them among themselves (such as Figure 1) please refer to supplementary material.

4.6. Statistics of isolated exemplars and miss-detections

We take a closer look into the discovered clusters. First, we try to understand the reason behind samples which end up in isolated exemplars. For this purpose we study the clusters with size less than 5. Then using the PASCAL annotations we compare their distribution (in terms of size and truncation) to the distribution of all training samples. It can be seen in figure 6 that the distribution of the isolated exemplars is biased toward less "unknown" samples. "unknown" means the sample is neither low resolution nor truncated.
Furthermore we study the statistics of the missed recalls on the PASCAL VOC **2008** dataset. Figure 7 depicts the recall-overlap plots for (1,2,5,10,100) highest scoring windows of each training image. Then we pinpoint the missed samples in the 100 highest scoring windows and show the statistics of those missed samples compared to all samples. Again it can be observed that the low resolution and truncated samples are the most problematic ones. It can be seen that we have a relatively very low number of missed samples which are neither low resolution, truncated nor occluded.
From the two figures it can be observed that the samples at tails of the distribution in our representation (HOG features) are mostly comprised of low resolution, occluded and truncated samples.

5. Conclusion and Future Works

In this work we have proposed a novel method of designing classifiers that tries to incrementally discover and learn visual subclasses of the semantic class. The formulation is such that it tries to find largest subset of the training samples that can be reliably distinguished from the negative class. The initial attempt has shown that the state of the art object detectors are unable to use around half of the data (on average) appearing at the tails of the distribution, thus showing potentials for future works. In fact we showed the performance is slightly improved by taking the outliers (isolated exemplars) out of the training set.
Results by ignoring the poor localization false positives in-

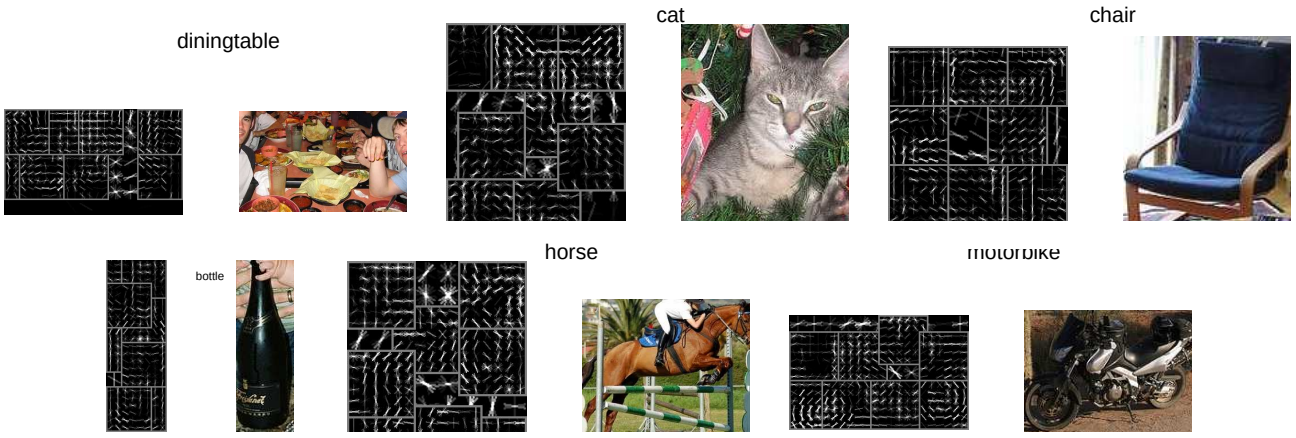

Figure 4: Trained filters for the largest visual subclass found for a subset of classes along with one of their members.

dicates that a suitable aggregation procedure for the output of the independently trained experts appears to be a promising future work to reduce false positives due to multiple counting. However, we think the most important future direction is the analysis of a way to enlarge or merge the discovered subclasses by including samples from the tails of the distribution. This can be done by re-designing the model, feature representations or by using additional data to fill in the identified gaps of the input space (between isolated exemplars).

6. Appendix A.

6.1. Individual subclass discovery and learning

We are interested in finding the *largest subset* of our positive training data ($y = 1$) which maintains a regular structure (large margin in case of L2 regularization) to *all* the negative data ($y = -1$). In the follwoing we first explain two alternative regularized risk minimization methods of largest subclass discovery. These methods use different approaches to calculate the induced loss and are inspired by self-paced learning and regularized ramp loss learners respectively.

$$\underset{\mathbf{w},\mathbf{v}}{\text{minimize}}\ \dot{J}(\mathbf{w}) := \Omega(\mathbf{w}) + \lambda \dot{R}(\mathbf{w}, \mathbf{v}, D) \tag{10}$$

6.2. Self paced subclass learning

Inspired by the self-paced learning framework of [22], we pose this problem as the following optimization,[3]

$$\dot{R}(\mathbf{w}, \mathbf{v}, D) = \sum_{i=1}^{N_p} v_i (l_i^{\mathbf{w}} - 1/K_p) + \sum_{i=1}^{N_n} v_{i+N_p} (l_{i+N_p}^{\mathbf{w}} - 1/K_n) \tag{11}$$

$$\mathbf{v} = \{v_1, v_2, ..., v_N\}, v_i \in \{0, 1\}$$

[3]For the simplicity of the notation, we assume samples are ordered by their labels thus in the set D there are N_p positive ($y_i = 1$) samples indexed by $i = 1 : N_p$ and then N_n negative ($y_i = -1$) samples indexed by $i = N_p + 1 : N_p + N_n$.

where $\mathbf{v}$ is a binary indicator vector that denotes which samples are considered in the learning of $\mathbf{w}$ parameters. Proportional to λ, K_p and K_n are scalars which control how predictable a sample should be to get in the subset. The larger the K the lower the loss of a sample should be in order to influence the learning of $\mathbf{w}$. The inclusion of more low loss samples ($l_i^w \leq 1/K$) is encouraged since it will decrease the total risk. Thus, K provides a trade-off between the inclusion of samples in learning and the amount of irregularity in $\mathbf{w}$ needed for that. This is in line with our intended trade-off of finding the largest subset while maintaining a level of regularity. Large K is a tighter constraint on the inclusion of the samples based on their closeness to the rest of the data.

This formulation is similar to Kumar et. al self-paced learning [22] where $K_p = K_n$. [22] proposed this framework for the non-convex latent variable models and in order to help them to converge to better local optima.

Optimizing 11 directly, involves integer programming for variable $\mathbf{v}$. However, for a convex choice of l and Ω the relaxed version (i.e. $v_i \in \mathbb{R}, 0 \leq v_i \leq 1$) is a biconvex programming in $\mathbf{w}$ and $\mathbf{v}$ whose optimal value can be shown to be at $v_i \in \{0, 1\}$. This means that when we fix either of the subset of variables $\mathbf{w}$ or $\mathbf{v}$ the problem becomes convex over the other subset. Biconvex problems are specially interesting since a simple alternating method has shown to converge to relatively good local optima efficiently [15].

6.3. Ramp loss expert learning

Finding the largest subset with high margin from negative samples can be seen as an outlier exclusion problem. That is, we want to detect those samples that are not coherent with (outliers w.r.t.) the largest majority of the data and exclude them from learning. This can be achieved by augmenting the loss function as follows,

$$\dot{R}(\mathbf{w}, D) = \sum_{i=1}^{N_p} min(l_i^{\mathbf{w}}, S_p) + \sum_{i=1}^{N_n} min(l_{i+N_p}^{\mathbf{w}}, S_n) \quad (12)$$

where the loss induced by each sample is capped at S_p for positive samples and S_n for negative samples. This means that those samples whose loss exceed the corresponding threshold will be discounted from optimization. One should note that, given specific model parameters $\mathbf{w}$, the gradient of high loss samples ($l_i^w > S$) is zero. Therefore, those samples does not influence the total risk by the change of $\mathbf{w}$ in any direction. It can be shown that the solution $\mathbf{w}^*$ of the above formulation is the same as if we have trained the model only on the samples with their $l_i^w <= S$. S in combination with λ controls the trade-off between regularization and inclusion of more samples. A large enough S will include all samples and as we decrease it towards it will be more strict to include samples.

This formulation is a general case of ramp loss SVM [7] where we have L2 regularization and hinge loss. [7] proposed ramp loss SVM to decrease the number of support vectors. This was in order to address scalability issues associated with learning non-linear SVMs.

It should be noted that for a convex choice of regularization Ω and sample loss functions 12 can be reformulated as difference of convex functions,

$$\dot{R}(\mathbf{w}, D) = \underbrace{\sum_{i=1}^{N_p} l_i^{\mathbf{w}} + \sum_{i=1}^{N_n} l_{i+N_p}^{\mathbf{w}}}_{\text{convex}} - \underbrace{[\sum_{i=1}^{N_p} max(l_i^{\mathbf{w}} - S_p, 0) + \sum_{i=1}^{N_n} max(l_{i+N_p}^{\mathbf{w}} - S_n, 0)]}_{\text{convex}} \quad (13)$$

This is true since ls are convex by choice, S_p and S_n are constants and convexity is preserved under max and summation operations. (13) can be solved using concave convex procedure (CCCP) [34].

6.4. Equivalence show case

It can be shown that the solutions to these two minimizations are equivalent. Here we give an intuition. Consider the case that we only have a single parameter $S = S_p = S_n = 1/K_p = 1/Kn$. Furthermore, take hinge loss as $l_i^{\mathbf{w}}$. In self-paced equation (11), for a given $\mathbf{w}$, we have the optimal values of indicator variables $v_i = 1$ iff $l_i^{\mathbf{w}} < S$ and $v_i = 0$ otherwise. That means, each sample i induces a self-paced loss,

$$v_i(l_i^{\mathbf{w}} - S) = \begin{cases} max(1 - y_i s_i, 0) - S & \forall i\, l_i^{\mathbf{w}} < S \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

On the other hand each samples induced ramp loss is $min(max(1 - y_i s_i, 0), S)$ by definition [7]. With a closer look at these two one see that self-paced induced loss is a shifted version of ramp loss by $-S$ and thus their solution are equivalent (look at figure 5).

6.5. Equivalence

It can be shown that the solutions to these two minimization are equivalent. Let's assume the $\mathbf{w}_r^*$ is the solution to the minimization (5). Without loss of generality, let's assume we have one parameter $S = S_p = S_n = 1/K_p = 1/Kn$. That gives,

$$J_{ramp}^* = \Omega(\mathbf{w}_r^*) + \sum_{i=1}^{N} min(l_i^{\mathbf{w}_r^*}, S) \leq \Omega(w) + \sum_{i=1}^{N} min(l_i^{\mathbf{w}}, S) \quad \forall \mathbf{w} \quad (15)$$

Now, let's assume that $\mathbb{N}_i^{\mathbf{w}}$ is the set of inlier samples that have an induced loss of $l_i^{\mathbf{w}} \leq S$ w.r.t. $\mathbf{w}$. And $\mathbb{N}_o^{\mathbf{w}}$ is the set of remaining samples that are outliers ($|\mathbb{N}_i^{\mathbf{w}}| + |\mathbb{N}_o^{\mathbf{w}}| = N$). Then from 15 we have,

$$J_{ramp}^* = \Omega(\mathbf{w}_r^*) + \sum_{i \in \mathbb{N}_i^{\mathbf{w}_r^*}} l_i^{\mathbf{w}_r^*} + |\mathbb{N}_o^{\mathbf{w}_r^*}|S \leq \Omega(\mathbf{w}) + \sum_{i \in \mathbb{N}_i^{\mathbf{w}}} l_i^{\mathbf{w}} + |\mathbb{N}_o^{\mathbf{w}}|S \quad \forall \mathbf{w} \quad (16)$$

On the other hand, let's assume $\mathbf{w}_s^*$ is one optimal solution to minimization (4). Optimal values of v_i will be one for all inliers and 0 for outliers as otherwise it would not be an optimal point. We also have $S = 1/K_p = 1/K_n$. Thus,

$$J_{self}^* = \Omega(\mathbf{w}_s^*) + \sum_{i \in \mathbb{N}_i^{\mathbf{w}_s^*}} l_i^{\mathbf{w}_s^*} - |\mathbb{N}_i^{\mathbf{w}_s^*}|S \leq \Omega(\mathbf{w}) + \sum_{i \in \mathbb{N}_i^{\mathbf{w}}} l_i^{\mathbf{w}} - |\mathbb{N}_i^{\mathbf{w}}|S \quad \forall \mathbf{w}$$

$$\overset{|\mathbb{N}_i^{\mathbf{w}}|+|\mathbb{N}_o^{\mathbf{w}}|=N}{\Longrightarrow} \Omega(\mathbf{w}_s^*) + \sum_{i \in \mathbb{N}_i^{\mathbf{w}_s^*}} l_i^{\mathbf{w}_s^*} + |\mathbb{N}_o^{\mathbf{w}_s^*}|S \leq \Omega(\mathbf{w}) + \sum_{i \in \mathbb{N}_i^{\mathbf{w}}} l_i^{\mathbf{w}} + |\mathbb{N}_o^{\mathbf{w}}|S \quad \forall \mathbf{w} \quad (17)$$

Comparing (16) to (17) shows that $\mathbf{w}_s^*$ and $\mathbf{w}_r^*$ are solutions to both minimizations (4) and (5).

6.6. Discussion

We want the visual subclasses to be well separated from *all* the negative samples. So we are interested in the special case of the above formulations when $S_n >> 0$ for ramp loss formulation and negative indicator variables set to 1 ($v_i = 1$) for self-tuned formulation. This will ensure that all the negative samples are included in the learning. Thus, we assume there is only one parameter K_p and S_p which we simple call K and S and in general *inclusion parameter*.

A larger choice of inclusion parameter will be a tighter constraint on the positive samples which are contributing to the final solution. That means by choosing large enough or small enough K or S in the formulations above we span a spectrum where we have exemplar learner [23] at one extreme and a monolithic learner at the other extreme respectively.

The two methods use different alternate searches to generate upper bound formulations (linear tangent in case of CCCP for ramp loss difference of convex formulation and fixed indicator variable for self-paced formulation). However, it can be shown that the generated bounds are closely related.

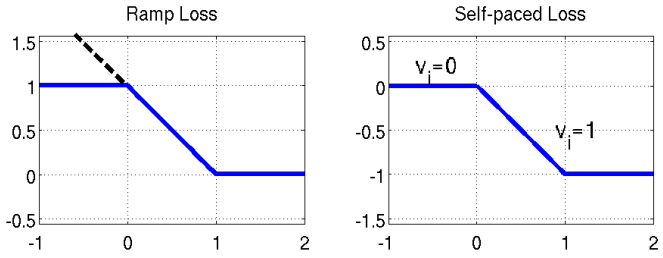

Figure 5: The loss that sample i induces with $S = 1/K = 1$ for the case of Ramp Loss (left) and Self-paced Loss (right). Dashed line shows the original hinge loss.

It can be easily shown that the slight difference is the tendency of the ramp loss approximated bound to keep outliers having low scores (remaining outlier) while self-paced formulation completely ignores the outliers in its approximated bound.

References

[1] O. Aghazadeh, H. Azizpour, J. Sullivan, and S. Carlsson. Mixture component identification and learning for visual recognition. In *ECCV*, pages 115–128, 2012. 2, 3, 7

[2] H. Azizpour and I. Laptev. Object detection using strongly-supervised deformable part models. In *ECCV*, pages 836–849, 2012. 3

[3] F. Bach and Z. Harchaoui. Diffrac: a discriminative and flexible framework for clustering. In *NIPS*, 2007. 3

[4] L. Bourdev and J. Malik. Poselets: Body part detectors trained using 3D human pose annotations. In *ICCV*, 2009. 3

[5] Y. Chen, L. Zhu, and A. Yuille. Active mask hierarchies for object detection. In *ECCV*, 2010. 6, 7

[6] R. Collobert, S. Bengio, and Y. Bengio. A parallel mixture of svms for very large scale problems. In *NIPS*, pages 633–640, 2001. 3

[7] R. Collobert, F. H. Sinz, J. Weston, and L. Bottou. Trading convexity for scalability. In *ICML*, pages 201–208, 2006. 5, 9

[8] C. Doersch, A. Gupta, and A. A. Efros. Mid-level visual element discovery as discriminative mode seeking. In *NIPS*, pages 494–502, 2013. 3

[9] C. Doersch, S. Singh, A. Gupta, J. Sivic, and A. A. Efros. What makes paris look like paris? *ACM Trans. Graph.*, 31(4):101, 2012. 3

[10] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. 1, 3

[11] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *PAMI*, 32(9):1627–1645, 2010. 2, 3, 4, 5

[12] R. B. Girshick, P. F. Felzenszwalb, and D. McAllester. Discriminatively trained deformable part models, release 5. http://people.cs.uchicago.edu/ rbg/latent-release5/. 6, 7

[13] R. Gomes, A. Krause, and P. Perona. Discriminative clustering by regularized information maximization. In *NIPS*, pages 775–783, 2010. 3

[14] R. Gomes, M. Welling, and P. Perona. Incremental learning of nonparametric bayesian mixture models. In *CVPR*, 2008. 3

[15] J. Gorski, F. Pfeuffer, and K. Klamroth. Biconvex sets and optimization with biconvex functions: a survey and extensions. *Math. Meth. of OR*, 66(3):373–407, 2007. 5, 9

[16] C. Gu, P. A. Arbeláez, Y. Lin, K. Yu, and J. Malik. Multi-component models for object detection. In *ECCV*, pages 445–458, 2012. 3, 6, 7

[17] C. Gu and X. Ren. Discriminative mixture-of-templates for viewpoint classification. In *ECCV*, pages 408–421, 2010. 3

[18] M. Hoai and A. Zisserman. Discriminative sub-categorization. In *CVPR*, pages 1666–1673, 2013. 3

[19] D. Hoiem, Y. Chodpathumwan, and Q. Dai. Diagnosing error in object detectors. In *ECCV (3)*, pages 340–353, 2012. 6, 7

[20] A. Kembhavi, B. Siddiquie, R. Miezianko, S. McCloskey, and L. S. Davis. Incremental multiple kernel learning for object recognition. In *ICCV*, pages 638–645, 2009. 3

[21] T.-K. Kim and R. Cipolla. Mcboost: Multiple classifier boosting for perceptual co-clustering of images and visual features. In *NIPS*, pages 841–856, 2008. 3

[22] M. P. Kumar, B. Packer, and D. Koller. Self-paced learning for latent variable models. In *NIPS*, pages 1189–1197, 2010. 3, 4, 5, 6, 8

[23] T. Malisiewicz, A. Gupta, and A. A. Efros. Ensemble of exemplar-svms for object detection and beyond. In *ICCV*, pages 89–96, 2011. 5, 7, 10

[24] A. J. O'Toole, P. J. Phillips, F. Jiang, J. H. Ayyad, N. Penard, and H. Abdi. Face recognition algorithms surpass humans matching faces over changes in illumination. *PAMI*, 29(9):1642–1646, 2007. 1

[25] B. S. P. Dollar, C. Wojek and P. Perona. Pedestrian detection: A benchmark. In *CVPR*, June 2009. 1

[26] N. Razavi, J. Gall, P. Kohli, and L. J. V. Gool. Latent hough transform for object detection. In *ECCV*, pages 312–325, 2012. 3

[27] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *CVPR DeepVision workshop*, 2014. 2

[28] X. Song, T. Wu, Y. Jia, and S.-C. Zhu. Discriminatively trained and-or tree models for object detection. In *CVPR*, 2013. 3

[29] A. Torsello, S. R. Bulò, and M. Pelillo. Beyond partitions: Allowing overlapping groups in pairwise clustering. In *ICPR*, pages 1–4, 2008. 3

[30] A. Vedaldi, V. Gulshan, M. Varma, and A. Zisserman. Multiple kernels for object detection. In *ICCV*, pages 606–613, 2009. 6, 7

[31] J. Wang and V. Saligrama. Local supervised learning through space partitioning. In *NIPS*, pages 91–99, 2012. 3

[32] Y. Yang and D. Ramanan. Articulated human detection with flexible mixtures of parts. *PAMI*, 35(12):2878–2890, 2013. 2

[33] T. Yeh and T. Darrell. Dynamic visual category learning. In *CVPR*, 2008. 3

[34] A. L. Yuille and A. Rangarajan. The concave-convex procedure. *Neural Computation*, 15:915–936, 2003. 9

[35] X. Zhu, C. Vondrick, D. Ramanan, and C. C. Fowlkes. Do we need more training data or better models for object detection? In *BMVC*, 2012. 1

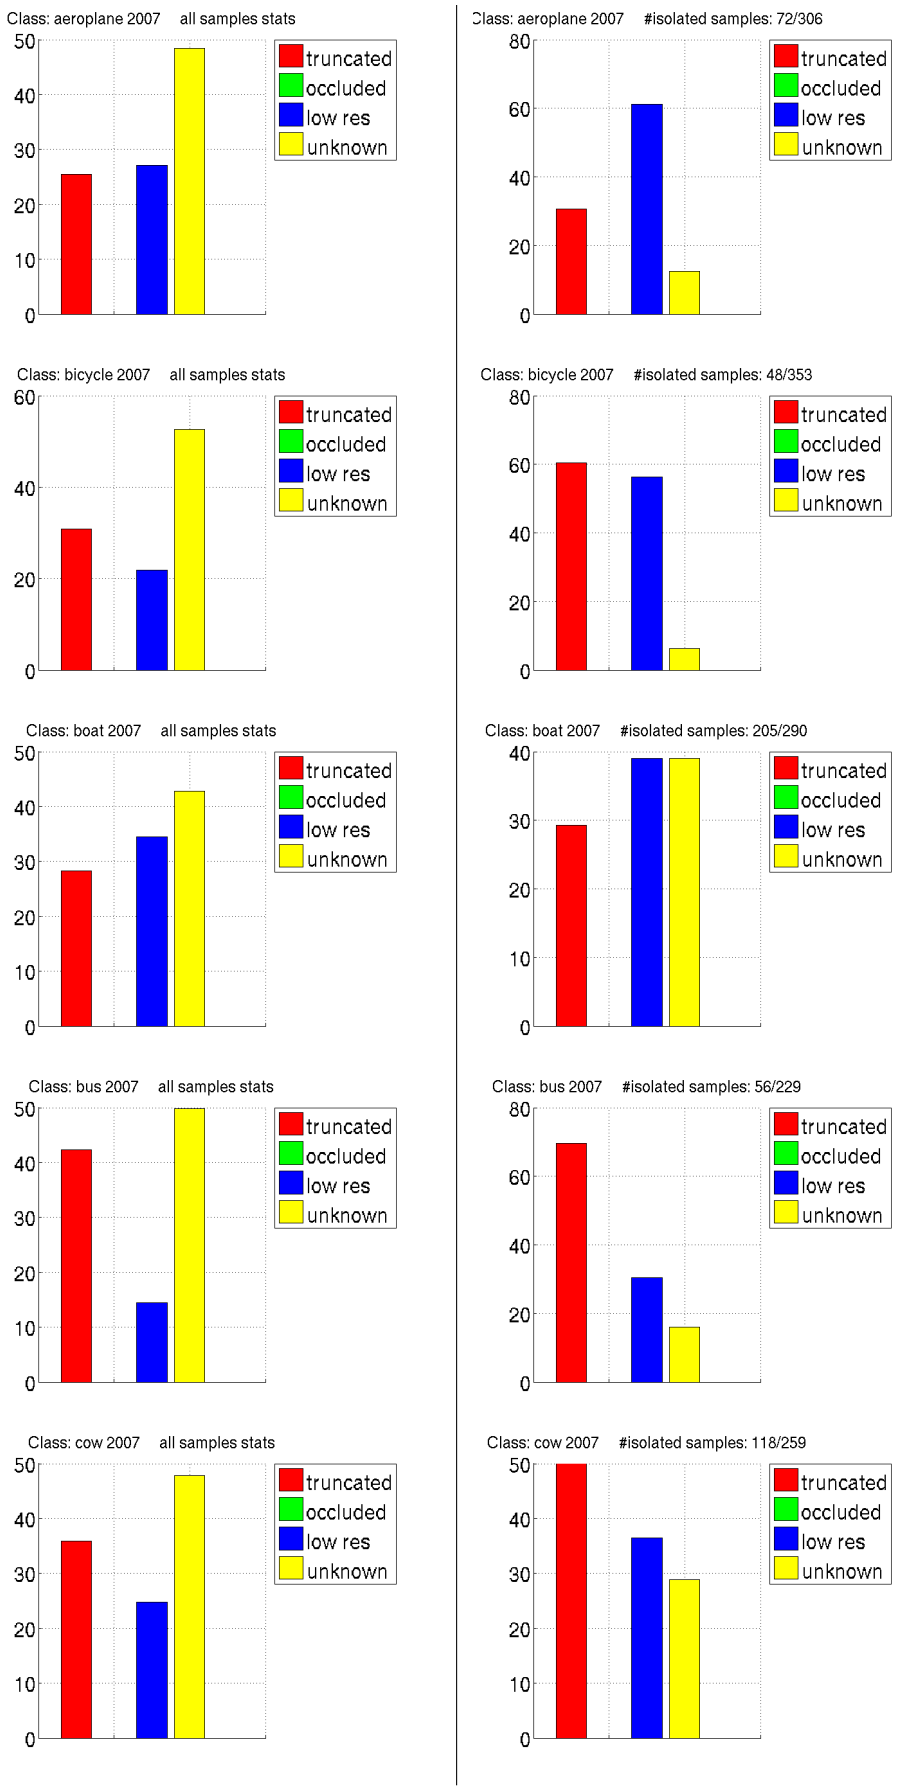
Class: aeroplane 2007 all samples stats
Class: aeroplane 2007 #isolated samples: 72/306
Class: bicycle 2007 all samples stats
Class: bicycle 2007 #isolated samples: 48/353
Class: boat 2007 all samples stats
Class: boat 2007 #isolated samples: 205/290
Class: bus 2007 all samples stats
Class: bus 2007 #isolated samples: 56/229
Class: cow 2007 all samples stats
Class: cow 2007 #isolated samples: 118/259
truncated
occluded
low res
unknown

Figure 6: Statistics on the distribution of isolated samples (members of clusters with size less than 5). The distribution is shown on three criteria: low resolution, truncated, unknown

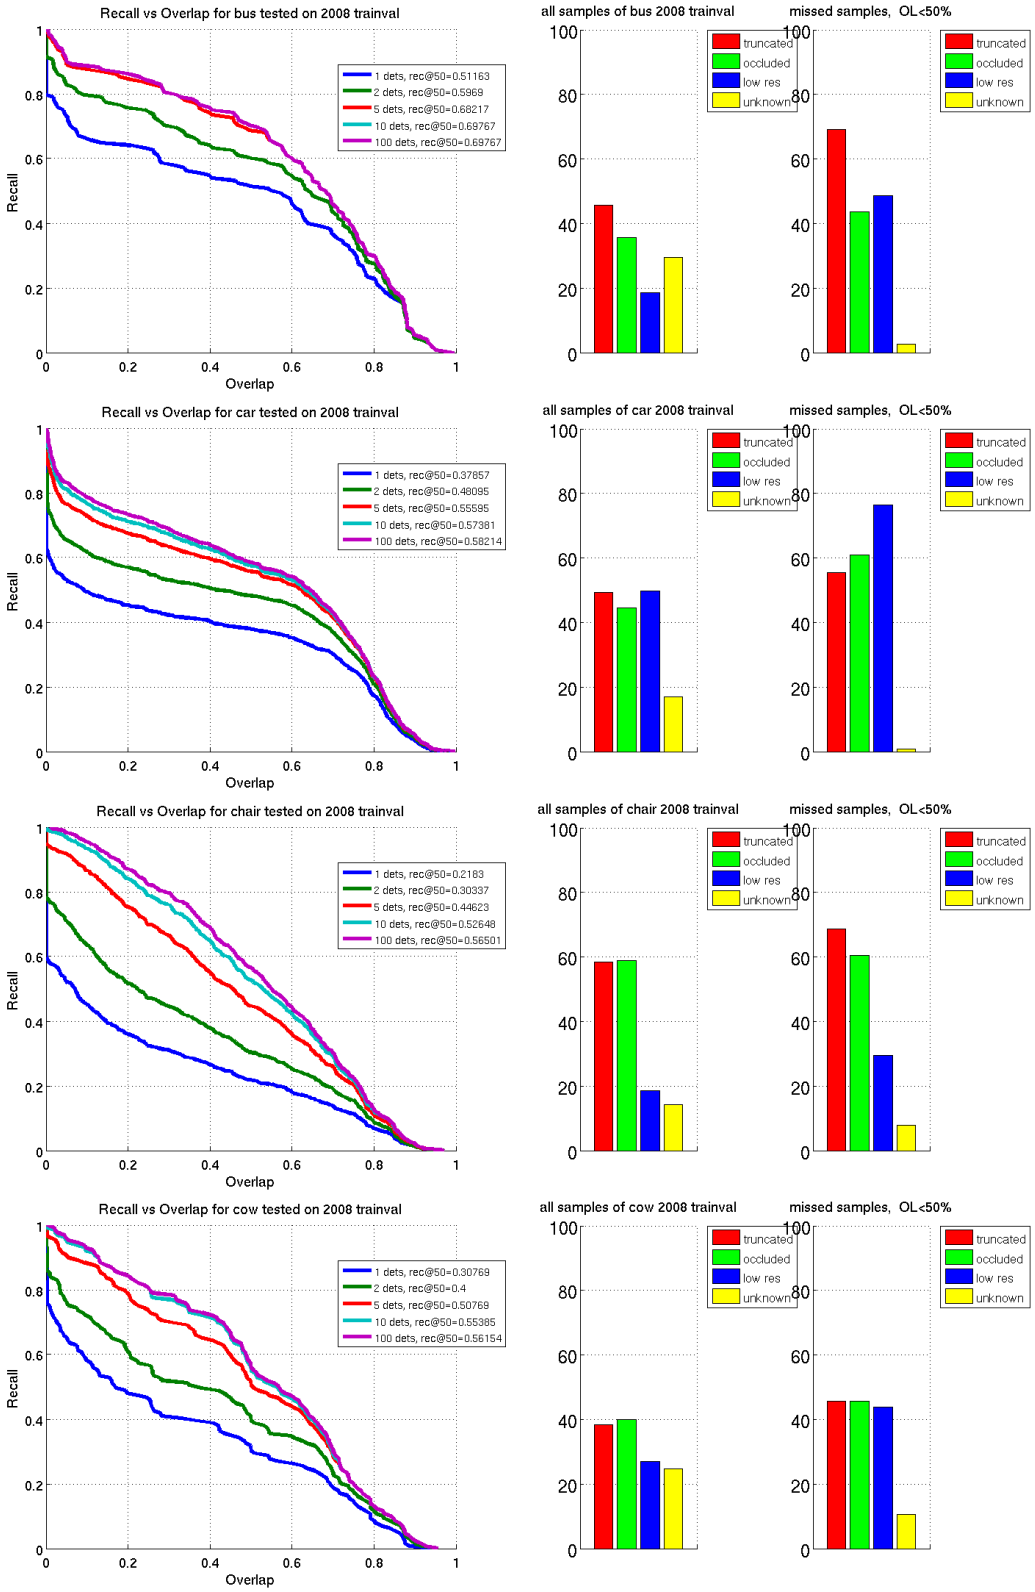

Figure 7: Statistics on the distribution of missed detection on PASCAL VOC 2008 training set. The distribution is shown on four criteria: low resolution, truncated, occluded, unknown. It can be seen that we have a relatively very low number of missed samples which are neither low resolution, truncated nor occluded.

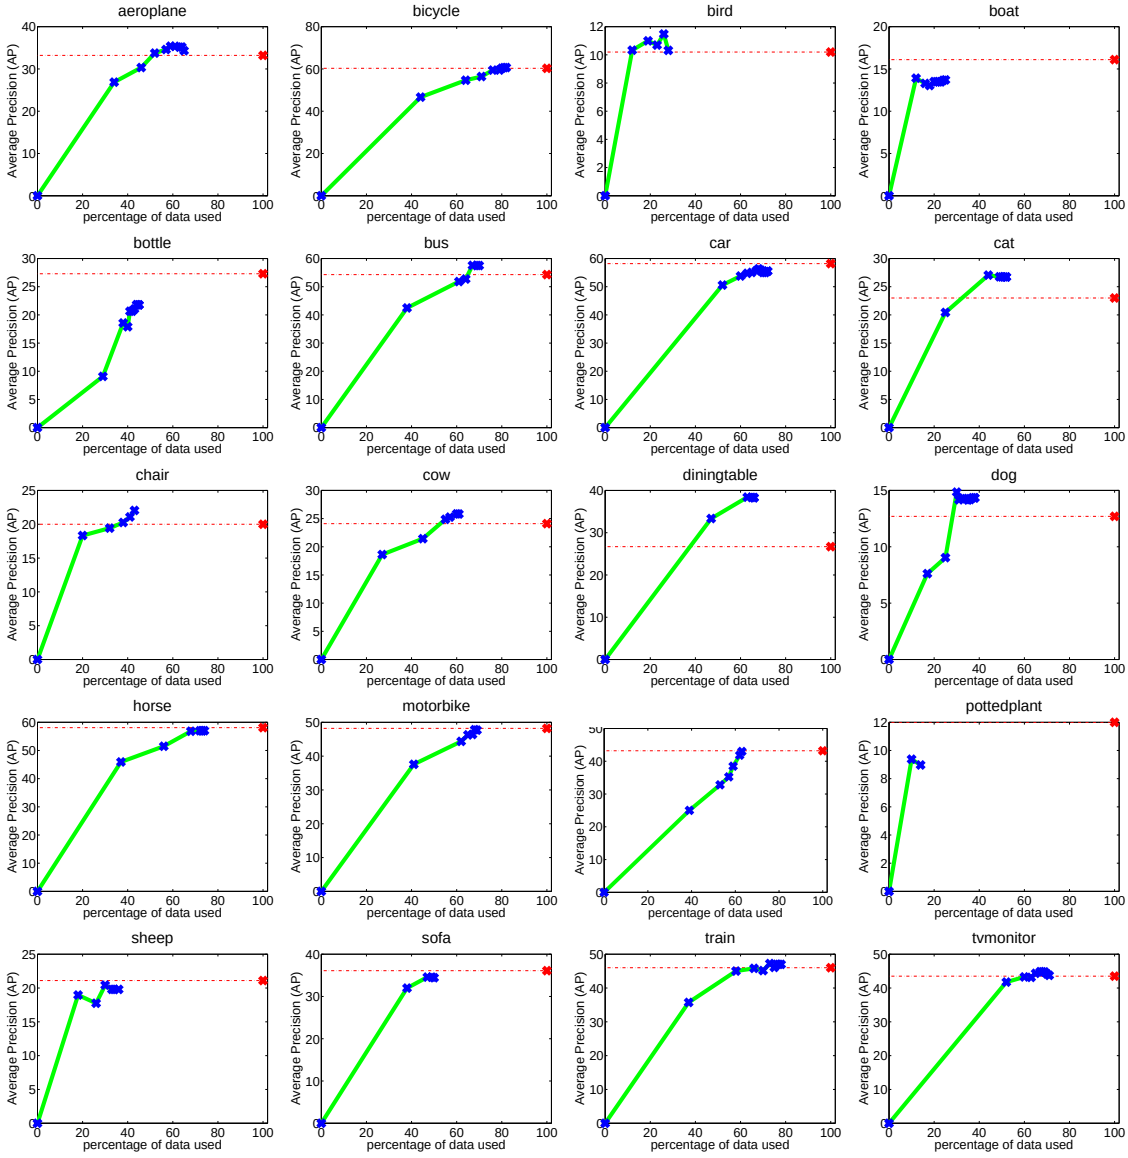

Figure 8: Evolution of the AP with more experts as a function of the fraction of the data used. Blue crosses in green (light gray) curve mark our method's AP after aggregation of each expert. Red cross (dashed line) is the AP of DPM using all data.